RECOVERY RATES, DEFAULT PROBABILITIES AND THE CREDIT CYCLE

Max Bruche and Carlos González-Aguado

CEMFI Working Paper No. 0612

September 2006

CEMFI
Casado del Alisal 5; 28014 Madrid
Tel. (34) 914 290 551. Fax (34) 914 291 056
Internet: www.cemfi.es

We would like to thank Ed Altman for giving us access to the Altman-NYU Salomon Center Default Database. We thank Rafael Repullo and Javier Suárez for helpful comments.

CEMFI Working Paper 0612
September 2006

RECOVERY RATES, DEFAULT PROBABILITIES AND THE CREDIT CYCLE

Abstract

Default probabilities and recovery rate densities are not constant over the credit cycle; yet many models assume that they are. This paper proposes and estimates a model in which these two variables depend on an unobserved credit cycle, modelled by a two-state Markov chain. The proposed model is shown to produce a better fit to observed recoveries than a standard static approach. The model indicates that ignoring the dynamic nature of credit risk could lead to a severe underestimation of e.g. the 95% VaR, such that the actual VaR could be higher by a factor of up to 1.7. Also, the model indicates that the credit cycle is related to but distinct from the business cycle as e.g. determined by the NBER, which might explain why previous studies have found the power of macroeconomic variables in explaining default probabilities and recoveries to be low.

JEL Codes: G21, G28, G33.
Keywords: Credit, recovery rate, default probability, business cycle, capital requirements, Markov chain.

Max Bruche
CEMFI
bruche@cemfi.es

Carlos González-Aguado
CEMFI
cgonzalez@cemfi.es

1 Introduction

In many first-generation commercial credit risk models, the default rate (the percentage of defaulting firms in each period) and recovery rates (the fraction of the face value of a bond that is recovered in case there is a default on the bond issue) are assumed to be independent. This is the case e.g. in J.P. Morgan's CreditMetrics™ (Gupton, Finger, and Bhatia, 1997) model. Alternatively, the models treat the recovery rate as a constant parameter, e.g. in CSFB's CreditRisk+™ (1997) model.

There is evidence, however, that periods of high recoveries are associated with low default rates and low recovery periods are associated with high default rates (Altman, Brady, Resti, and Sironi, 2005; Frye, 2000, 2003; Hu and Perraudin, 2002). From a pricing as well as a risk management perspective, it therefore seems unreasonable to assume that default rates are independent of recovery rates.

This has led to a second generation of credit risk management models that take into account the covariation of default rates and recovery rates, e.g. Standard & Poor's LossStats™and KMV Moody's LossCalc™(Gupton and Stein, 2002).

Often, it is informally argued that both variables reflect the state of economy, i.e. the business cycle.[1] This paper formalizes a very similar idea, which is that both variables depend on the *credit cycle*. We propose a model that uses only the default rate and recovery rate data to draw conclusions about the state of the credit cycle, which is treated as an otherwise unobserved variable,[2] whose dynamics are described by a two-state Markov chain.

This model is tested in various ways, and shown to be a reasonable representation of the data. Integrating information about industry and seniority, we show how the model can be used to calculate loss distributions for portfolios. Allowing for dependence between default probabilities and recovery rates via the state of the credit cycle can increase e.g. the 95% VaR of a given portfolio by a factor of up to 1.7, and even in credit upturns, the VaR can increase by a factor of up to 1.3. Altman, Brady, Resti, and Sironi (2005) obtain a very similar difference in VaRs on the basis of a hypothetical simulation exercise; we can confirm on the basis of our estimated model that their numbers are

[1] Additionally, Acharya, Bharath, and Srinivasan (2005) find that the condition of the *industry* is important.

[2] In spirit, this is closest to the analysis of Frye (2000), although he proceeds in the context of an essentially static model based on Vasicek (1987).

very, very plausible.

The difference in VaRs is economically significant, and should be taken into account e.g. by banks attempting to implement the Advanced Internal Ratings-Based Approach of Basel II, under which they can calculate their own default probabilities and estimates of loss given default. We propose that banks use unconditional probabilities of being in a credit upturn or downturn for calculating capital requirements in the context of a dynamic model such as the one we present. This would take into account the dynamic nature of credit risk, as well as ensuring that capital requirements are not cyclical.

We compare our credit downturns to recession periods as determined by the NBER. The begin of our credit downturns typically precede the start of a recession, and continue until after the end of a recession (see e.g. figure 2). This indicates that the credit cycle is related to, but distinct from the macroeconomic cycle. We argue that this explains why previous studies have found that macroeconomic variables explain only a small proportion in the variation of recovery rates (Altman, Brady, Resti, and Sironi, 2005).

The increase in risk associated with negatively correlated default frequencies and recovery rates is likely to have an effect on pricing, and might (at least in part) explain the "corporate spread puzzle" (see e.g. Chen, Collin-Dufresne, and Goldstein, 2006). This will be explored in future versions of this paper.

The rest of this paper is structured as follows: The model is presented in section 2. In section 3 we describe the data set used. Section 4 discusses the estimation and various tests, and section 5 explores the implications for credit risk management. Finally, section 6 concludes.

2 The model

We allow default probabilities and recovery rates to be related by letting them both depend on the state of the credit cycle. Conditional on the cycle, default probabilities and recovery rates are independent. The state of the cycle will be assumed to be described by an unobserved two-state Markov chain. In each state we have a different default probability, i.e. there will be a parameter describing the default probability in an upturn, and a parameter describing the default probability in a downturn. The number of defaults will be binomially distributed according to these probabilities. For each default, a recovery rate will be drawn from a density that depends on the state, and also on issuer characteristics (industry) and bond characteristics (seniority). Sometimes, firms default on issues with different seniorities simultaneously, so we also need to specify how recoveries on these issues are then drawn simultaneously.

The parametric density chosen for the marginal density of recovery rates is the beta density, which is often used by practitioners (see e.g. Gupton and Stein, 2002). Beta distributions are well suited to modelling recoveries, as they have support $[0,1]$, and in spite of requiring only two parameters (α and β) are quite flexible.

As a consequence of having probabilities of being in an upturn or downturn of the credit cycle, the particular density describing possible recoveries at any particular point in time will be a mixture of the two different beta densities for the two different states, with the mixing probabilities being given by the probabilities of being in either state. This gives considerably more flexibility in matching observed unconditional distributions of recovery rates, and is in some sense similar to the nonparametric kernel density approach of Renault and Scaillet (2004).[3]

An advantage of making parametric assumptions is that estimation via maximum likelihood procedures is possible. Here, we proceed by employing a slightly modified version of the method proposed by Hamilton (1989).

2.1 The likelihood function

Let d_t be the number of defaulted firms observed in t. Firms will be indexed by i. Each firm might have issues of different seniority classes, which we index by c (c ranges from 1 to C).

Let Y_t be a matrix of dimension $d_t \times C$ containing observations on the d_t defaulted firms, one row per firm, for all different seniority classes c. Denote the typical element as y_{tic}, and denote each row in this matrix as $y_{ti.}$.

We treat situations in which one firm defaults on several different issues, possibly with different seniorities, within a given period indexed by t as a single default event.

Then our objective is to maximize the following likelihood function

$$L = \sum_{t=1}^{T} \log f(Y_t, d_t | \Omega_{t-1}), \tag{1}$$

[3]It is similar in the sense that we are also using a mixture of beta distributions. The difference is that obviously, a kernel density estimator will mix over a much larger number of densities (equal to the number of observations), with mixing probabilities constant and equal across the kernel densities.

where Ω_{t-1} is all information available at $t-1$, which includes observed defaults up to that point. Note that we have made explicit that there is information in both the matrix of recoveries Y_t and the number of its rows, d_t, about the state, i.e. identification of the state is obtained through both the number of defaults and recoveries.

Letting s denote state ($s = 1$ denotes credit upturn and $s = 0$ denotes credit downturn), we can then write

$$f(Y_t, d_t|\Omega_{t-1}) = f(Y_t, d_t|s_t = 1, \Omega_{t-1}) \Pr(s_t = 1|\Omega_{t-1}) + f(Y_t, d_t|s_t = 0, \Omega_{t-1}) \Pr(s_t = 0|\Omega_{t-1}). \quad (2)$$

Assume that the values of the recovery rates in the rows of Y_t, $y_{ti.}$, are *independent* of d_t conditional on the state. Once the state is known, the number of defaults does not contain information about the likely values of recoveries, and the values of recoveries do not contain information about the likely number of defaults. This assumption is of course crucial, and its validity will be examined below (see section 4).

The assumption of conditional independence allows us to write

$$f(Y_t, d_t|\Omega_{t-1}) = f(Y_t|s_t = 1, \Omega_{t-1}) \Pr(d_t|s_t = 1, \Omega_{t-1}) \Pr(s_t = 1|\Omega_{t-1}) + f(Y_t|s_t = 0, \Omega_{t-1}) \Pr(d_t|s_t = 0, \Omega_{t-1}) \Pr(s_t = 0|\Omega_{t-1}). \quad (3)$$

The contribution to the likelihood function of a given period is given by the sum of two products, one for each state. The components of these products are the state-conditional density function of the observed recovery rates, the state-conditional probability of the number of defaults and the probability of being in that state. These components are now described separately.

2.1.1 State-conditional recoveries

The density function of the observed recovery rates can be simplified if we assume that (conditional on the state) recoveries of different firms of the same period are independent draws from the same distribution. This means that recoveries on different default events (as defined above) are independent. Again, this conditional independence assumption allows us to factorize, so the conditional densities can now be written as

$$f(Y_t|s_t = j, \Omega_{t-1}) = \prod_{i=1}^{d_t} f(y_{ti.}|s_t = j, \Omega_{t-1}) \quad (4)$$

Both of these conditional independence assumptions (independence of d_t and the value of recoveries, and independence of recoveries of different firms) imply that the dependence between the number of defaults, and the recoveries, and the dependence between defaults of two firms is entirely driven by the state of the business cycle.

We allow the parameters of our marginal distributions of recovery rates to vary according to state j, seniority class c and industry k. The marginal density is assumed to be given by

$$f(y_{tic}|s_t = j, \Omega_{t-1}) = \frac{1}{B(\alpha_{jck}, \beta_{jck})} y_{tic}^{\alpha_{jck}-1}(1 - y_{tic})^{\beta_{jck}-1} \tag{5}$$

where the parameters vary according to state j, seniority class c and industry k. This does not address the dependence structure across seniority classes (for the same default event), however.

Clearly, recoveries across different seniority classes but for the same default event cannot be assumed to be independent. It is likely that a high firm-level recovery implies higher instrument-level recoveries, and that therefore recoveries on instruments issued by the same firm but with different seniorities exhibit positive dependence. Also, it would be natural to expect that higher seniority classes observe a higher recovery. As it turns out, in the data, it is relatively frequently the case that higher seniority classes recover less than the lower seniority classes.

In the absence of a model that relates firm-level recoveries to instrument-level recoveries (see e.g. Carey and Gordy, 2004, for a discussions of some of the issues), we propose the relatively simple assumption that the dependence structure of recovery rates across different seniority classes for the same firm is given by a Gaussian copula, such that we can specify a correlation matrix of recoveries across seniority classes. We do not impose a dependence structure that implies that higher seniority will always recover more, since this is not what we observe in the data.[4]

2.1.2 State-conditional number of defaults

The probability of observing d_t defaults is binomial (conditional on the state). Let N_t be the number of (defaulted and non-defaulted) firms observed in period t and r_j the

[4]We considered conditioning recoveries of junior classes on the observed recoveries in senior classes via different ways of truncating densities. This did not lead to a better fit of the model, and in general complicated testing.

default probability in state j. Then we can calculate the probability of observing d_t in state j as

$$\Pr(d_t|s_t = j, \Omega_{t-1}) = \begin{pmatrix} N_t \\ d_t \end{pmatrix} r_j^{N_t}(1-r_j)^{N_t-d_t} \tag{6}$$

2.1.3 Probabilities of the states

Finally, the probability of being in one state can be derived from the total probability theorem. Obviously,

$$\begin{aligned} \Pr(s_t = 1|\Omega_{t-1}) = \quad & \Pr(s_t = 1|s_{t-1} = 1, \Omega_{t-1}) \Pr(s_{t-1} = 1|\Omega_{t-1}) \\ & + \Pr(s_t = 1|s_{t-1} = 0, \Omega_{t-1}) \Pr(s_{t-1} = 0|\Omega_{t-1}). \end{aligned}$$

To shorten notation, let p denote the transition probability of state 1 to state 1 ($\Pr(s_t = 1|s_{t-1} = 1, \Omega_{t-1})$), and q denote the transition probability of state 0 to state 0 (($\Pr(s_t = 0|s_{t-1} = 0, \Omega_{t-1})$)), which will be the two parameters in our Markov chain, we can write

$$\begin{aligned} \Pr(s_t = 1|\Omega_{t-1}) = \quad & p \Pr(s_{t-1} = 1|\Omega_{t-1}) + (1-q) \Pr(s_{t-1} = 0|\Omega_{t-1}) \\ = \quad & (1-q) + (p+q-1) \Pr(s_{t-1} = 1|\Omega_{t-1}) \end{aligned} \tag{7}$$

The probability $\Pr(s_{t-1} = 1|\Omega_{t-1})$ can be obtained via a recursive application of Bayes' rule:

$$\begin{aligned} & \Pr(s_{t-1} = 1|\Omega_{t-1}) = \\ & \qquad \frac{f(Y_{t-1}|s_{t-1} = 1, \Omega_{t-2}) \Pr(d_{t-1}|s_{t-1} = 1, \Omega_{t-2}) \Pr(s_{t-1} = 1|\Omega_{t-2})}{f(Y_{t-1}|\Omega_{t-2})} \end{aligned} \tag{8}$$

Note that this specification is consistent with the existence of periods without observations (periods of time in which there are no defaults and hence no recoveries are observed). In this case, the conditional densities in [3] will disappear and the contribution to the likelihood function of this period will take into account only the number of defaults, not the recovery rates:

$$\begin{aligned} f(Y_t|\Omega_{t-1}) = & \Pr(s_t = 1|\Omega_{t-1}) \Pr(d_t|s_t = 1, \Omega_{t-1}) \\ & + \Pr(s_t = 0|\Omega_{t-1}) \Pr(d_t|s_t = 0, \Omega_{t-1}) \end{aligned} \tag{9}$$

3 Data

The data is extracted from the Altman-NYU Salomon Center Corporate Bond Default Master Database. This data set consists of more than 2,000 defaulted bonds of US firms from 1974 to 2005. Each entry in the database lists the name of the issuer of the bond (this means that we can determine its industry as described by its SIC code), the seniority of the bond, the date of default and the price of this bond per 100 dollars of face value one month after the default event.

Typically, the database contains the prices of many bonds for a given firm and seniority class on a given date, so we need to aggregate. We do this by taking weighted averages (weighted by issue size).[5]

We also aggregate data across time into periods corresponding to calendar years. We assume that a default of the same firm within twelve months of an initial default event (e.g. in December and then March of the following year) represent a single default event.

We calculate the recovery rate as the post-default price divided by the face value. Some of the recovery rates calculated in this way are larger than 1. This probably reflects the value of coupons. We scale the recovery rates by a factor of .9 to ensure that our observations lie in the support of the beta distribution.

The data set does not contain the total size of the population of firms from which the defaults are drawn. This variable is necessary in order to calculate the probability of observing d_t defaults in [6]. To determine the population size we use default frequencies as reported in Standard & Poor's Quarterly Default Update from May 2006 (taken from the CreditPro Database). Dividing the number of defaulting firms in each year in the Altman data by Standard & Poor's default frequency, we can obtain a number for the total population of firms under the assumption that both data sets track the same set of firms. The available default frequency data ranges from 1981 to 2005, so we lose observations from periods 1974-1980. However, these years cover only 9 observations in the Altman data.

After these adjustments, the final data set contains 1,078 observations. Descriptive statistics for these observations are presented in Tables 1 to 3.

Table 1 reports the yearly statistics of our adjusted data set. The first column

[5] This is something commonly used in the literature. See for example Varma and Cantor (2005).

shows the annual default frequencies as reported by Standard & Poor's, while the other columns refer to the main data set, containing the number of observations per year and their means and standard deviation. It can already be seen that typical recession years (as published by the NBER: 1990-91 and 2001) show higher default frequencies and lower recoveries than years of economic expansion. This will of course play an important role below.

Tables 2 and 5 report the number of observations and the mean and standard deviation of recoveries classified by seniority and industry respectively. These are in line with those reported in other papers, although on average recoveries are slightly lower here. In Table 2, it is shown that mean recovery increases with increasing seniority, while standard deviation remains more or less constant across seniorities. In Table 5, we see that the mean recovery is highest for utilities and lowest for telecoms.

4 Estimation and tests

We present estimates for several versions of the model. Ultimately, we are interested in constructing a model that takes into account the effects of the credit cycle, industry and seniority. Initially, however, we examine how well the model does using only the assumption about state dependence, ignoring the effects of industry and seniority (i.e. we look at a version of the model where marginal recovery rate distributions do not vary according to industry and seniority). Also, to sidestep the issue of dependence of recoveries on issues of the same firm, but with different seniorities, we initially drop all default events for which we observe recoveries on more than one seniority (the remaining observations are roughly representative of the whole sample). We estimate a *basic static model*, in which recovery rate distributions and default probabilities do not depend on the state, and contrast this with a *basic dynamic model* in which both recovery rate distributions and default probabilities are state dependent.

We then proceed towards a model that takes into account industry and seniority. Since we do not have a sufficient number of default events for every possible combination of industry, seniority and state of the credit cycle, however, we will need to aggregate industries into broad groups.

In order to check that the aggregation is reasonable, we first allow marginal distributions to vary across industries only, and check which industries can be grouped.

Next, we allow marginal distributions to vary across seniorities only. We can then look at the dependence of recoveries on issues of the same firm but with different seniorities. For this we need to add the observations dropped at the earlier stage back in, and look at the full data set.

Finally, we combine information on industry and seniority to construct the *dynamic industry/ seniority model*. This will be contrasted with a *static industry/ seniority model*.

At the various stages, different tests will be performed.

4.1 Basic model

4.1.1 The basic static model

In the *static model* recovery rate distributions and default probabilities do not depend on the state. We also ignore industry and seniority. We estimate on the sub-sample of observations for which we do not observe recoveries across different seniorities for the same firm.

Estimates of this model are provided in Table 6. With the estimated distribution parameters we obtain an implied mean recovery of 37%.[6] The default probability in this static case is found by taking the average of all the default frequencies in our data set and it is equal to 1.47%.

4.1.2 The basic dynamic model

In the *basic dynamic model* recovery rate distributions and default probabilities depend on the state, but we ignore industry and seniority. We estimate on the sub-sample of observations for which we do not observe recoveries across different seniorities for the same firm.

The estimated parameters are reported in tables 6 and 7. Average recoveries are much lower in downturns (31% versus 47%), and default probabilities are higher (2.69% versus 0.86%).

We test whether the estimated parameters values for downturns are significantly different from the estimated parameter values for upturns via likelihood ratio tests: We

[6] The mean of a beta distribution with parameters α and β is $\frac{\alpha}{\alpha+\beta}$. Note that we also have to divide the resulting number by 0.9 in order to undo the scaling adjustment.

first test whether the recovery rate distribution parameters parameters are different, and then test whether the default probabilities are different. The p-value are less than 0.01% for both tests, so that the null hypothesis of no difference across states is rejected both for recovery rate distributions and default probabilities. The credit cycle matters.

We can also compare the integrated squared error (ISE) between more flexible beta kernel density estimates (Renault and Scaillet, 2004) and our dynamic and static recovery rate density estimates respectively. To make this comparison straightforward, we pick periods for which the smoothed probabilities of being in a downturn or an upturn are unambiguous (i.e. either 1 or 0 respectively). These ISEs are presented in Table 8. It can be seen that that the ISE for the static model is between 5 and 11 times larger, indicating that the beta distributions indicated by the static model are much further away from the empirical density (the beta kernel density estimate) than the densities estimated on the basis of our dynamic model.

A more complete method for evaluating density forecasts was proposed by Diebold, Gunther, and Tay (1998), described in the appendix, section A. The basic idea is that applying the probability integral transform based on the predicted (filtered) distributions to the actual observations (the recovery rates in our case) should yield an i.i.d.-uniform series (the PIT series) under the null hypothesis that the density forecasts are correct. Departures from uniformity are easily visible when plotting histograms, and departures from i.i.d. are visible when plotting autocorrelation functions of the PIT series.

Examining the histograms and correlograms of the transformed series (see figures 3 to 6) indicates that both models seem to get the marginal distribution wrong (as indicated by the histogram), but that the dynamic model is much better at explaining the dynamics of the recovery rate series (as indicated by the correlogram). It is likely that the high positive autocorrelation visible in the PIT series of the static model is due to the fact that recoveries depend on the state of the credit cycle, which is persistent. This is of course ignored by the static model.[7]

We also plot the (smoothed) probabilities of being in the credit downturn over time and compare these to NBER recession dates in figure 1. As can be seen, the credit cycle

[7]We also test that the unconditional distribution of the PIT series is uniform via a Kolmogorov-Smirnov test, and obtain p-values of 2.32% for the basic static model, and 6.38% for the basic dynamic model respectively. Note that these p-values are strictly speaking incorrect because they ignore parameter uncertainty.

had two major downturns, around the recession of the early 90s and around 2001, but in each case, the credit downturn started well before the recession and ended after it. To emphasize this point, we also estimate the model on quarterly data. The resulting plot (figure 2) shows this difference between recessions and credit downturns more clearly.[8] This difference between the credit cycle and the business cycle might explain the low explanatory variable of macroeconomic variables documented by Altman, Brady, Resti, and Sironi (2005).

We also examine our assumption that recoveries and default probabilities are independent conditional on the state of the credit cycle (see equation 3). We regress annual mean recoveries on default rates for the whole sample. The coefficient is significantly different from zero at 5% (the correlation between the two variables is -0.43). We also run two regressions separately in the two subsamples given by the periods for which the smoothed probabilities of being in a credit upturn or downturn are unambiguous (i.e. either 1 or 0). The coefficients of the regressions in the two sub-samples are not significantly different from zero at 5% (the correlations are 0.47 in upturns and -0.12 in downturns). This seems to support our assumption.

4.2 Adding information on industry

We now let the parameters of the recovery rate density depend on industry, as well as on the unobserved credit cycle. We have 12 industries, which implies that we need to estimated 48 recovery rate distribution parameters (2 states, 12 industries, 2 parameters per beta distribution), 2 transition probabilities and 2 default probabilities. The estimated parameters of the recovery rate distributions are provided in Table 10. The implied mean recoveries are reported in Table 11. Transition and default probability estimates are virtually identical to the estimates in the basic model and therefore not reported.

As discussed, in order to construct a combined industry / seniority model, due to data limitations, we will have to aggregate. We aggregate industry into three groups:

1. Group A: Financials, Leisure, Transportation, Utilities

2. Group B: Consumer, Energy, Manufacturing, Others

[8] The "semi-downturn" of 1986 consists mostly of defaulting oil companies. It is probable that this is related to the drop in oil prices at the time.

3. Group C: Building, Mining, Services, Telecoms

For the given sample, these groupings roughly correspond to industries with high, medium and low mean recovery rates respectively. We test whether the parameters of the recovery rate distribution are significantly different within these groups (i.e. whether it the groupings are reasonable). The p-value of the likelihood ratio test is 0.10, which indicates that the null of the same parameters within the groups cannot be rejected at conventional levels of significance.

4.3 Adding information on seniority

We now estimate a model in which the parameters of the recovery rate densities depend on the state of the credit cycle and seniority (and ignore industry for the time being). We observe 4 different categories of seniority (Senior Secured, Senior Unsecured, Senior Subordinated, Subordinated). The estimated recovery rate parameters are shown in Table 13. The mean recovery rates implied by these parameters are reported in Table 14. We observe that although senior secured bonds have a higher recovery in upturns on average (52%), they have very similar (low) recoveries to bonds of all other seniorities in downturns (29%). As has been pointed out by Frye (2000), this might have important consequences for risk management, as instruments that are though of as "safe" turn out to be safe only in good times.

Again, transition and default probability estimates are virtually identical to those in the basic model and therefore not reported.

4.3.1 Dependence across seniorities

So far we have dealt only with observations for which we do not observe several recoveries across different seniorities for the same firm, which we now add back in. It is unreasonable to assume that these observations of recoveries represent independent draws. As described above, letting our marginal recovery rate densities be beta densities as before, we model this dependence with a Gaussian copula.

In order to estimate this dependence, we add the observations of default events for which we observe recovery rates across instruments of different seniorities back in. For these observations, seniority seems to mean something quite different, as can be seen from comparing Table 3 and 4. One explanation for this would be that recovery on a bond that is labelled as “senior unsecured” (for example) depends not so much on this label, but much more on whether or not there exists a cushion of junior debt.

Since categories of seniority seem to have a very different meaning when we observe recoveries across instruments of different seniorities, we allow the parameters of our marginal recovery rate densities for a given seniority to be distinct for the cases where either recovery only on a single seniority is observed, or whether recovery is observed on more than one seniority.

The estimated parameters of the marginal beta densities are reported in Table 15, and the estimated correlation matrix is presented in Table 16.

Implied mean recoveries of this model are presented in Table 17. We can see that junior debt recovers less and senior debt recovers more for default events for which we observe recoveries for more than seniority. For example, in an upturn, Senior Unsecured debt would recover (on average) 46% and Subordinated debt would recover 34% if recoveries are observed on more than one seniority, whereas they would recover (on average) 42% and 37% respectively if recovery was only observed on a single seniority.

At 5% significance, only the correlations between Senior Unsecured, Senior Subordinated and Subordinated seniorities are significantly different from zero. These correlations are positive, indicating that a high recovery on e.g. Senior Unsecured debt for a defaulting firm indicates a likely high recovery on its Senior Subordinated Debt. The recovery on Senior Secured debt seems to be less strongly related to recoveries on different seniorities. This could reflect that the existence of a junior debt cushion matters less, and the quality of the collateral matters more for Senior Secured debt.

4.4 The industry / seniority dynamic model

Having decided on groupings of industries such that we now have a sufficient amount of data for each possible combination of seniority group, industry group, and state of the business cycle, we are now in a position to estimate a combined industry / seniority model. The recovery rate density parameters are in tables 18 and 19, again the transition and default probability parameters are virtually identical to the ones in the basic model.

For comparison purposes, we also estimate a static industry / seniority dynamic model, which is identical with the dynamic industry / seniority model except for the fact that the static model ignores the credit cycle. We calculate the PIT series for both models, and plot histograms and correlograms in figures 7, 8, 9, and 10. We can see that the unconditional distributions seem to be closer to uniformity, suggesting that taking into account industry and seniority helps in describing distributions of recovery rates. Looking at the correlograms, it is also apparent, however, that the static version of the industry / seniority model is again unable to explain the dynamics of recovery rates. We conclude that allowing the parameters of the density of recoveries to depend on the credit cycle allows for a better match of the empirically observed recovery rates.[9]

5 Implications for risk management

Altman, Brady, Resti, and Sironi (2005) suggest that a correlation between default rates and recovery rates can imply much higher VaR numbers than those implied by independence between these two variables. In a completely static model, they compare VaRs calculated in the case of independence, and in the case of perfect rank correlation between default rates and recovery rates.

We are in a position to calculate the loss distributions and statistics of loss distributions (such as the VaR) implied by our estimated model; i.e. taking into account the estimated degree of dependence between default probabilities and recovery rates, which in our case is driven entirely by the credit cycle.

We calculate (by simulation of 10,000 paths) the one-year loss distribution of a hypothetical portfolio of 500 bonds. For this calculation, we have to chose the probability that we attach to being in a credit downturn today. We examine cases with this probability being equal to one (we know that we are in a downturn today), zero (we know that we are in an upturn today) and 33.5%, which corresponds to the unconditional probability of being in a downturn, given our estimated transition probabilities. This means that we are calculating loss distributions from the point of view of a manager that has no information about the state of the credit cycle.

Comparing the *basic static model* and the *basic dynamic model* (i.e. ignoring industry and seniority for the time being), we can see from Table 20 that the 95% VaR is a 2.39% loss on the portfolio assuming the world is dynamic as described in our model, versus a 1.58% loss on the portfolio if we had assumed that the world is static. This is

[9] We also calculate a Kolmogorov-Smirnov test on the PIT series, and find a p-value of 9.91% for the static model, and a p-value of 14.37% for the dynamic model. Note that the caveat mentioned above applies.

a very sizeable difference in risk. It arises because the dynamic model allows for credit downturns, in which not only the default rate is very high, but also, the recovery rate is very low. This amplifies losses vis-a-vis the static case.

Even supposing that we are in an upturn today, the dynamic model still produces a 95% VaR of 1.96%, which is larger than the static VaR because even though we are in an upturn today, we might go into a credit downturn tomorrow, with higher default rates and lower recoveries.

The loss density implied by the dynamic model based on the unconditional probability of being in a downturn is compared to the loss density implied by the static model in figure 11. The dynamic model loss density is bimodal, reflecting the possibility of ending up either in an upturn with low default probabilities and high recoveries, or ending up in a downturn, with high default probabilities and low recoveries. As can be seen, the tail of the loss distribution implied by the dynamic model is much larger.

We also show this comparison for the case of assuming that we are in an upturn today in figure 12. It can be seen that even being in an upturn today, the possibility of going into a downturn tomorrow produces a bimodality in the loss distribution (albeit smaller) implied by the dynamic model, and a larger tail than that produced by the static model.

Looking at our industry / seniority model, we can see that the underestimation of the VaR of the static model seems to be most pronounced for Group C, our low recovery industries (including e.g. Telecoms), where the VaR based on the dynamic model can be up to 1.7 times as large.

5.1 Procyclicality

Altman, Brady, Resti, and Sironi (2005) argue that estimates of loss given default and of the default probability should be countercyclical, which would lead to capital requirements calculated under the Basel II Internal Ratings-based Approach to be procyclical. Concretely, capital requirements should increase in credit downturns, as the estimated loss given default and the default probability rise simultaneously. They note that regulation should encourage banks to use "long-term average recovery rates" (and presumably default probabilities) in calculating capital requirements to dampen the procyclicality of capital requirements.

Looking at Table 20, we can see that indeed if banks were to use a dynamic model like the one presented here, capital requirements would be procyclical since the VaR is higher for credit downturns. Our take on encouraging banks to use "long-term average recovery rates" would be that banks should use a dynamic model, but use unconditional probabilities of being in downturns to calculate capital requirements. In this case, capital requirements would not be cyclical, but would still take into account the dynamic nature of risk.

6 Conclusions

This paper formalizes the idea that default probabilities and recovery rates depend on the credit cycle: it proposes a formal model that is estimated and found to fit the data reasonably well in various ways. We demonstrate that taking into account the dynamic nature of credit risk in this way implies higher risk, since default probabilities and recovery rates are negatively related. The estimated credit downturns seem to start much earlier and end later than the recessions as reported by the NBER, indicating that the credit cycle is to some extent distinct from the business cycle, which might explain why Altman, Brady, Resti, and Sironi (2005) find that macroeconomic variables in general have low explanatory power for recovery rates. The paper shows that if banks and regulators were to fully take into account the dynamic nature of credit risk this would imply the cyclicality of capital requirements, as has been argued elsewhere. We argue that using the estimated unconditional probabilities of being in credit downturn based on a dynamic model such as the one presented here would meet the twin goals of taking into account the dynamic nature of risk as well as ensuring that capital requirements are not cyclical. Lastly, we hinted that the larger probabilities of large losses implied by the dynamic model presented here might explain part of the corporate spread puzzle. This issue will be examined in future versions of this paper.

References

Acharya, V. V., S. T. Bharath, and A. Srinivasan, 2005, "Does Industry-wide Distress Affect Defaulted Firms? - Evidence from Creditor Recoveries," London Business School Working Paper.

Altman, E., B. Brady, A. Resti, and A. Sironi, 2005, "The Link between Default and Recovery Rates: Theory, Empirical Evidence, and Implications," *Journal of Business*, 78, 2203–2227.

Carey, M., and M. Gordy, 2004, "Measuring Systemic Risk in Recoveries on Defaulted Debt I: Firm-Level Ultimate LGDs," Working Paper.

Chen, L., P. Collin-Dufresne, and R. S. Goldstein, 2006, "On the Relation Between the Credit Spread Puzzle and the Equity Premium Puzzle," Working Paper.

Credit Suisse First Boston, 1997, "CreditRisk+: A Credit Risk Management Framework," New York.

Diebold, F. X., T. A. Gunther, and A. S. Tay, 1998, "Evaluating Density Forecasts with Applications to Financial Risk Management," *International Economic Review*, 39, 863–883.

Frye, J., 2000, "Depressing Recoveries," Risk Magazine.

Frye, J., 2003, "A False Sense of Security," Risk Magazine.

Gupton, G. M., C. C. Finger, and M. Bhatia, 1997, "CreditMetrics™- Technical Document," J.P. Morgan, New York.

Gupton, G. M., and R. M. Stein, 2002, "LossCalc™: Model for Predicting Loss Given Default (LGD)," Moody's KMV, New York.

Hamilton, J. D., 1989, "Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357–384.

Hu, Y.-T., and W. Perraudin, 2002, "Recovery Rates and Defaults," Birbeck College Working Paper.

Renault, O., and O. Scaillet, 2004, "On the way to recovery: A nonparametric bias free estimation of recovery rate densities," *Journal of Banking and Finance*, 28, 2915–2931.

Standard & Poor's, 2006, "Quarterly Default Update and Rating Transitions," New York.

Varma, P., and R. Cantor, 2005, "Determinants of Recovery Rates on Defaulted Bonds and Loans for North American Corporate Issuers: 1983-2003," *Journal of Fixed Income*, 14, 29–44.

Vasicek, O., 1987, "Probability of loss on loan portfolio," KMV Working Paper.

A The Diebold-Gunther-Tay test

Diebold, Gunther, and Tay (1998) propose a test (DGT test) for evaluating density forecasts. The basic idea is that under the null hypothesis that these forecasts are equal to the true densities (conditioned on past information), applying the cumulative distribution function to the series of observations should yield a series of iid uniform-$[0, 1]$ variables. Whether the transformed variables are iid uniform can be tested in various ways, e.g. via a Kolmogorov-Smirnov test. Departures from uniformity or iid are easily visible by looking at histograms and autocorrelation functions of the transformed series.

In order to apply the DGT test to our predicted recovery rate densities, we create a vector $y^{\dagger}$ with typical element $y_t^{\dagger} = \text{vec}\left(Y_t^{\mathsf{T}}\right)$. For each element in the vector, we can now create a density forecast from our estimated model which conditions on all previous elements in this vector, and uses the filtered state probabilities. Due to our independence assumptions, this is mostly straightforward (most previous elements do not enter the conditional distributions). A slight complication arises due to the assumption of a Gaussian copula between recoveries of the same firm but across different seniorities. For these observations, conditional densities can easily be calculated via the conditional Gaussian distribution, though.

Applying the cumulative distribution function associated with these density forecasts to the vector $y^{\dagger}$ yields a vector of transformed variables which we call z. Under the null hypothesis that the density forecasts are correct, the elements of z should be an iid uniform series. Serial correlation of the series would indicate that we have not correctly conditioned on the relevant information. A departure from uniformity would indicate that the marginal distributions are inappropriate.

B Tables and Figures

B.1 Descriptive Statistics Tables

Year	Default frequency	Number of observations	Mean Recovery	Standard Deviation
1981	0.14%	1	12.00	-
1982	1.18%	12	39.64	14.27
1983	0.75%	5	48.24	20.35
1984	0.90%	11	48.88	16.59
1985	1.10%	14	48.17	21.28
1986	1.71%	26	35.19	18.16
1987	0.94%	19	52.89	27.05
1988	1.42%	35	37.19	20.33
1989	1.67%	41	43.55	28.29
1990	2.71%	81	25.49	21.80
1991	3.26%	94	40.37	26.27
1992	1.37%	37	51.50	24.02
1993	0.55%	21	37.58	19.61
1994	0.61%	16	43.77	24.88
1995	1.01%	24	43.76	24.69
1996	0.49%	18	43.59	23.79
1997	0.62%	23	54.95	23.76
1998	1.31%	32	46.55	24.52
1999	2.15%	94	30.29	19.92
2000	2.36%	112	28.03	23.81
2001	3.78%	137	24.71	18.05
2002	3.60%	100	29.78	16.69
2003	1.92%	58	39.24	23.48
2004	0.73%	35	50.59	24.13
2005	0.55%	32	58.71	23.41

This table reports some annual statistics of the sample used in the paper. First column figures are default frequencies extracted from Standard & Poor's (2006). The other three columns are the number of observations and the mean and standard deviation for recovery rates.

Table 1: Recovery Rates Statistics by Year

Seniority	Number of observations	Mean Recovery	Standard Deviation
Senior Secured	210	42.26	25.76
Senior Unsecured	376	36.86	23.54
Senior Subordinated	334	32.73	23.66
Subordinated	158	34.17	23.00

This table reports the number of observations and the mean and standard deviation of recovery rates in our sample classified by seniority.

Table 2: Recovery rates by seniority

Seniority	Number of observations	Mean Recovery	Standard Deviation
Senior Secured	158	40.11	23.90
Senior Unsecured	276	35.62	22.09
Senior Subordinated	226	33.69	23.29
Subordinated	90	37.91	20.21

This table reports the number of observations and the mean and standard deviation of recovery rates in our sample classified by seniority, for default events with recovery on instruments of a single seniority only.

Table 3: Recovery rates by seniority (single seniority only)

Seniority	Number of observations	Mean Recovery	Standard Deviation
Senior Secured	52	48.81	29.77
Senior Unsecured	100	40.30	26.85
Senior Subordinated	108	30.70	24.28
Subordinated	68	29.24	25.42

This table reports the number of observations and the mean and standard deviation of recovery rates in our sample classified by seniority, for default events with recoveries accross more than one seniority only.

Table 4: Recovery rates by seniority (multiple seniorities only)

Industry	Number of observations	Mean Recovery	Standard Deviation
Building	15	32.19	30.32
Consumer	152	35.56	22.89
Energy	50	37.65	17.09
Financial	104	36.91	25.99
Leisure	53	46.03	27.77
Manufacturing	368	35.78	22.89
Mining	15	35.05	18.67
Services	76	33.41	25.83
Telecom	123	31.53	21.14
Transportation	66	38.99	24.01
Utility	23	46.93	28.14
Others	33	38.01	23.44

This table reports the number of observations and the mean and standard deviation of recovery rates in our sample classified by industry.

Table 5: Recovery rates by industry

B.2 Basic Model Tables

Static specification		Dynamic specification			
		Downturn		Upturn	
α	β	α_0	β_0	α_1	β_1
1.4474	2.9288	1.4181	3.5990	1.9860	2.7241

This is the basic model (no industry and seniority), in the static version (i.e. default probability and recovery rate distribution are constant, there is no credit cycle) and the dynamic version (default probability and recovery rate distribution depend on the credit cycle).

Table 6: Recovery rate density parameters (basic model)

	est. value	95% confidence interval
p	0.8707	0.6006 - 0.9679
q	0.7408	0.3645 - 0.9344
r_0	0.0269	0.0246 - 0.0295
r_1	0.0086	0.0076 - 0.0097

Table 7: Transition and default probabilities (basic model)

	Basic dynamic model	Basic static model
Downturn	0.0250	0.1370
Upturn	0.0204	0.2310

Only periods with probabilities of downturn / upturn of one are considered. ISE $= \int_0^1 (f_m(x) - f_e(x))^2 \, dx$, where f_m is the model density, and f_e is the empirical density (i.e. the beta kernel density estimate).

Table 8: Integrated Squared Error

	p-value
Basic static model	0.0232
Basic dynamic model	0.0638

Table 9: Kolmogorov-Smirnov test of DGT-transformed recoveries

B.3 Industry model tables

	α_0	β_0	α_1	β_1
Building	0.8783	6.0405	2.9565	4.2768
Consumer	1.6379	3.4976	1.7818	2.8330
Energy	4.8878	8.4901	3.6112	5.7415
Financial	1.2534	1.7697	1.7054	2.1217
Leisure	1.7519	2.9508	1.4490	1.8521
Manufacturing	1.4409	4.0063	2.2361	2.7325
Mining	2.7132	7.1598	2.9327	3.6264
Services	0.9950	3.0480	2.6556	4.1652
Telecom	1.6347	5.2445	2.7992	5.0656
Transportation	2.1564	5.3268	1.3672	2.4168
Utility	1.0017	2.2649	3.3697	2.7382
Others	2.1649	5.6149	3.5737	4.5783

Table 10: Recovery rate density parameters (industry model)

	Downturn	Upturn
Building	0.1269	0.4087
Consumer	0.3189	0.3861
Energy	0.3654	0.3861
Financial	0.4146	0.4456
Leisure	0.3725	0.4389
Manufacturing	0.2645	0.4500
Mining	0.2748	0.4471
Services	0.2461	0.3893
Telecom	0.2376	0.3559
Transportation	0.2882	0.3613
Utility	0.3067	0.5517
Others	0.2783	0.4384

Table 11: Mean recoveries (industry model)

Group A (High RR):	Financial, Leisure, Transp., Utility
Group B (Medium RR):	Consumer, Energy, Manuf., Others
Group C (Low RR):	Building, Mining, Services, Telecom

The LR test for assuming that recoveries within these groups are drawn from the same beta distributions has a p-value of 0.1003

Table 12: Groupings of industries

B.4 Seniority Model Tables

	α_0	β_0	α_1	β_1
Senior Secured	1.5080	3.7242	2.8179	2.6014
Senior Unsecured	1.5872	4.1010	2.1664	2.9563
Senior Subordinated	1.2165	3.1425	1.4265	2.2403
Subordinated	1.6246	3.4969	2.8277	4.7524

Table 13: Recovery rate density parameters (seniority model)

	Downturn	Upturn
Senior Secured	0.2903	0.5225
Senior Unsecured	0.2844	0.4275
Senior Subordinated	0.2828	0.3954
Subordinated	0.3196	0.3720

Table 14: Mean recoveries (seniority model)

B.4.1 Dependence across Seniorities

Single seniority only	α_0	β_0	α_1	β_1
Senior Secured	1.4953	3.6863	2.8207	2.6349
Senior Unsecured	1.5809	4.0832	2.1788	2.9797
Senior Subordinated	1.2117	3.1484	1.4170	2.2870
Subordinated	1.4951	3.1791	2.8452	4.9285
Multiple seniorities only	α_0	β_0	α_1	β_1
Senior Secured	1.3292	2.5514	5.8617	2.2875
Senior Unsecured	1.0724	2.4544	1.4327	1.6664
Senior Subordinated	0.8875	2.4601	1.1786	2.1021
Subordinated	0.6590	2.8774	1.5457	3.0064

Table 15: Recovery rate density parameters (Dependence across Seniorities Model)

	S.Sec.	S.Uns.	S.Sub.	Sub.
S.Sec.	1.0000			
S.Uns.	0.2942	1.0000		
S.Sub.	0.1935	0.5449	1.0000	
Sub.	-0.2427	0.4976	0.8008	1.0000

NOTE: None of the correlations involving Senior Secured are different from zero at a 5% level of significance.

Table 16: Correlations of recoveries across different seniorities of the same firm

Single seniority only	Downturn	Upturn
Senior Secured	0.2886	0.5170
Senior Unsecured	0.2791	0.4224
Senior Subordinated	0.2779	0.3826
Subordinated	0.3199	0.3660
Multiple seniorities only	Downturn	Upturn
Senior Secured	0.3425	0.7193
Senior Unsecured	0.3041	0.4623
Senior Subordinated	0.2651	0.3592
Subordinated	0.1864	0.3396

Table 17: Mean recoveries (Dependence across Seniorities Model)

B.5 Industry / Seniority Model

Single seniority only		α_0	β_0	α_1	β_1
Group A	Senior Secured	1.3481	2.8525	2.9140	2.5464
	Senior Unsecured	1.5668	2.5826	1.5688	2.9588
	Senior Subordinated	1.7561	3.7186	0.9822	1.2307
	Subordinated	1.0515	1.4742	2.6826	4.5252
Group B	Senior Secured	2.0018	5.5988	2.9346	2.8156
	Senior Unsecured	1.6957	4.2500	2.4316	2.8778
	Senior Subordinated	1.2955	3.1757	1.6778	3.0337
	Subordinated	2.1275	5.4169	2.5997	4.1187
Group C	Senior Secured	1.3022	3.3210	1.5760	2.1701
	Senior Unsecured	1.9317	7.0903	3.4871	5.4421
	Senior Subordinated	0.9438	3.0932	1.8059	2.7522
	Subordinated	1.3569	4.1573	2.5959	4.3902
Multiple seniorities only		α_0	β_0	α_1	β_1
Group A	Senior Secured	2.7358	2.8585	4.2663	1.9468
	Senior Unsecured	0.9257	1.6984	1.2148	2.6416
	Senior Subordinated	0.8376	2.4929	0.8468	1.3719
	Subordinated	0.6408	2.5614	1.5556	3.4942
Group B	Senior Secured	1.4238	4.5619	2.8220	1.7865
	Senior Unsecured	1.1351	1.6423	2.2624	2.5336
	Senior Subordinated	0.9744	2.7496	2.2301	4.7745
	Subordinated	0.6795	3.1701	2.1047	5.7251
Group C	Senior Secured	1.5581	3.4703	2.6236	1.2105
	Senior Unsecured	0.9945	3.0158	4.7929	3.0807
	Senior Subordinated	0.7572	1.3666	2.2040	2.5248
	Subordinated	0.5384	1.2774	2.8524	3.6745

Table 18: Recovery rate density parameters (Industry / Seniority Model)

	S.Sec.	S.Uns.	S.Sub.	Sub.
S.Sec.	1.0000			
S.Uns.	0.4216	1.0000		
S.Sub.	0.1606	0.5281	1.0000	
Sub.	-0.2429	0.6911	0.7595	1.0000

NOTE: None of the correlations involving Senior Secured are different from zero at a 5% level of significance.

Table 19: Correlations of recoveries across different seniorities of the same firm

B.6 Risk implications

		Static specification	Dynamic specification		
			Upturn	Unconditional	Downturn
Basic model		0.0158	0.0196	0.0239	0.0263
Group A	Senior Secured	0.0148	0.0194	0.0235	0.0260
	Senior Unsecured	0.0155	0.0176	0.0214	0.0240
	Senior Subordinated	0.0153	0.0189	0.0233	0.0258
	Subordinated	0.0153	0.0166	0.0203	0.0227
Group B	Senior Secured	0.0160	0.0209	0.0252	0.0280
	Senior Unsecured	0.0159	0.0200	0.0245	0.0273
	Senior Subordinated	0.0171	0.0196	0.0246	0.0272
	Subordinated	0.0159	0.0202	0.0246	0.0276
Group C	Senior Secured	0.0172	0.0203	0.0246	0.0275
	Senior Unsecured	0.0184	0.0221	0.0268	0.0299
	Senior Subordinated	0.0179	0.0215	0.0269	0.0294
	Subordinated	0.0169	0.0216	0.0265	0.0288

Table 20: 95% VaR for some selected models.

B.7 Figures

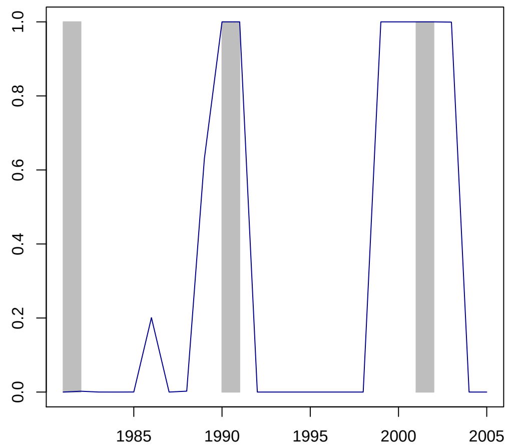

This figure plots the annual smoothed probability of being in a credit downturn as estimated on the basis of the basic dynamic model (no industry and seniority), assuming that the last quarter of 1981 was a credit downturn with a probability corresponding to the unconditional probability of being in a downturn. This is contrasted with NBER recessions (grey areas).

Figure 1: Probabilities of being in the credit downturn versus NBER recessions (annual)

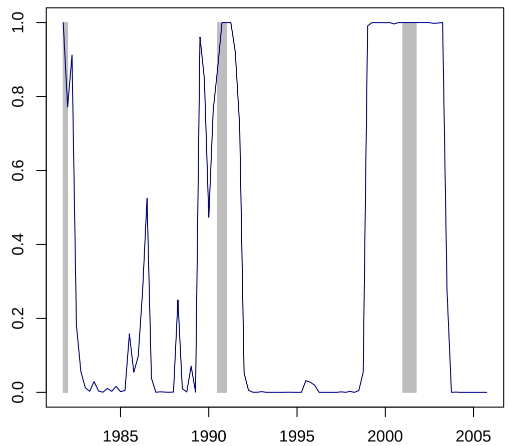

This figure plots the quarterly smoothed probability of being in a credit downturn as estimated on the basis of the basic dynamic model (no industry and seniority), assuming that the last quarter of 1981 was a credit downturn. This is contrasted with NBER recessions (grey areas).

Figure 2: Probabilities of being in the credit downturn versus NBER recessions (quarterly)

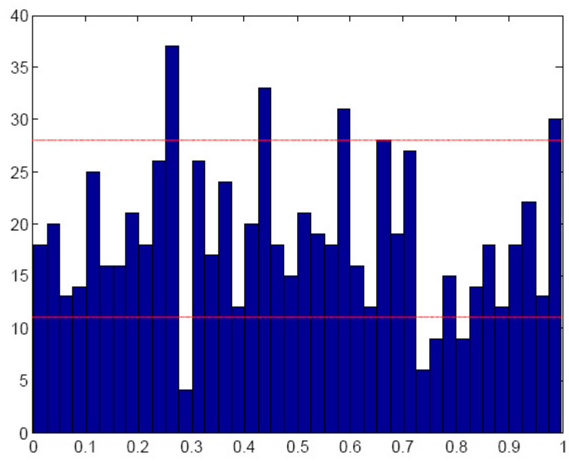

Figure 3: Histogram of the DGT-transformed recoveries - basic static model

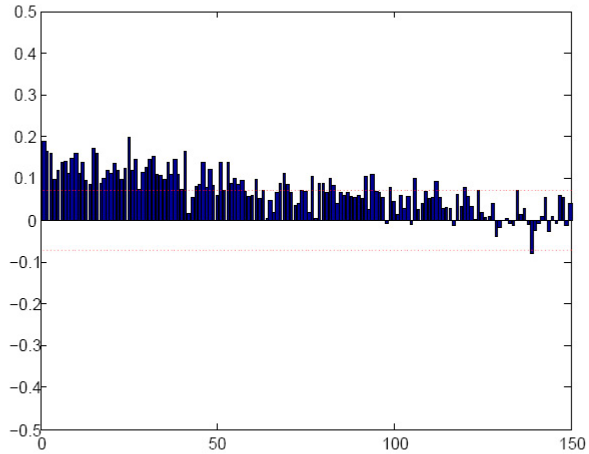

Figure 4: Correlogram of the DGT-transformed recoveries - basic static model

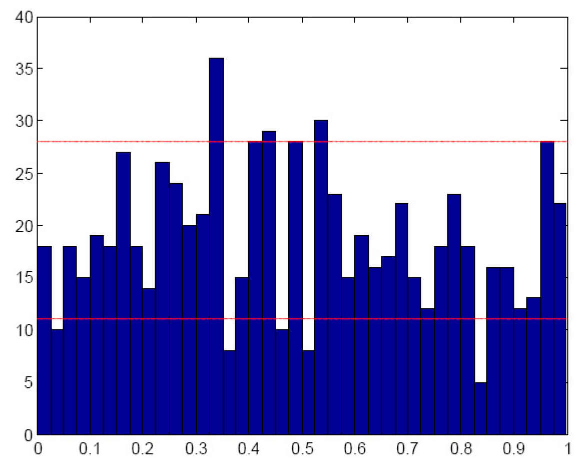

Figure 5: Histogram of the DGT-transformed recoveries - basic dynamic model

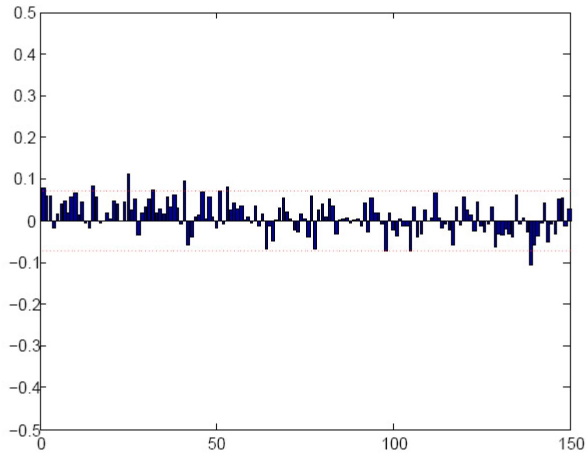

Figure 6: Correlogram of the DGT-transformed recoveries - basic dynamic model

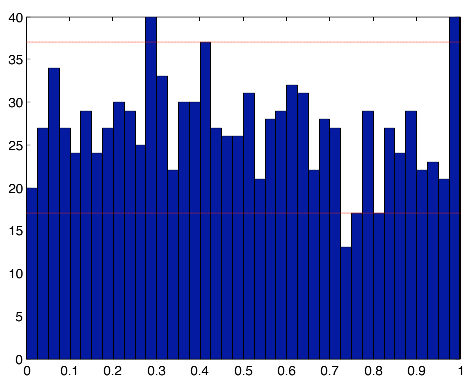

Figure 7: Histogram of the DGT-transformed recoveries - full static model

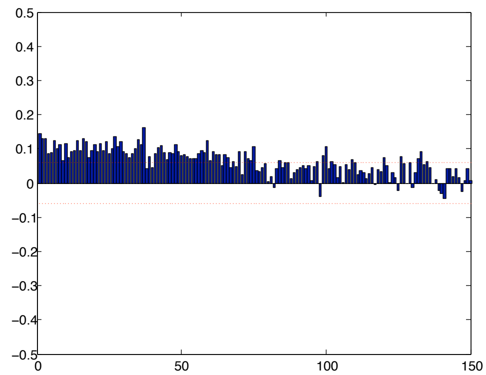

Figure 8: Correlogram of the DGT-transformed recoveries - full static model

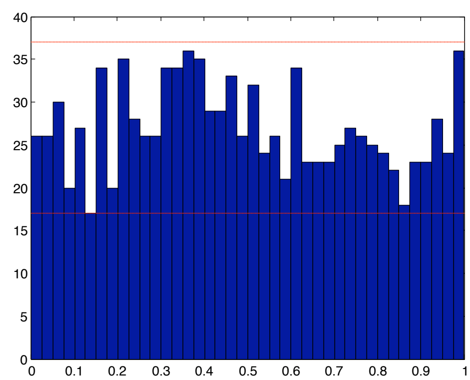

Figure 9: Histogram of the DGT-transformed recoveries - full dynamic model

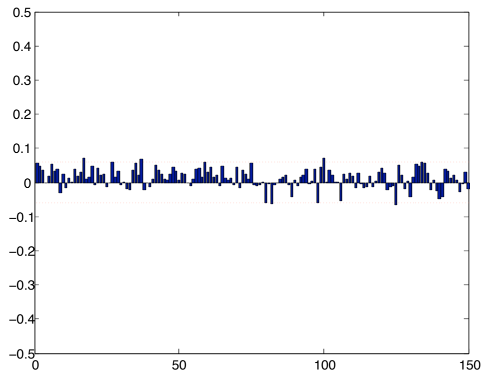

Figure 10: Correlogram of the DGT-transformed recoveries - full dynamic model

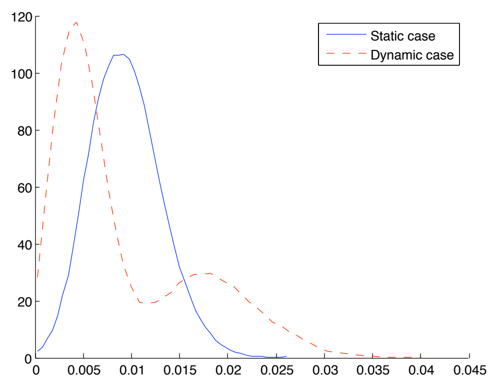

This figure contrasts the simulated loss density (pdf), of the basic model (no industry and seniority), for the static version (i.e. default probability and recovery rate distribution do not vary with the credit cycle) and the dynamic version (default probability and recovery rate distribution depend on the credit cycle). For the dynamic model, the probability of being in a credit downturn is assumed to be equal to the unconditional probability

Figure 11: Simulated loss density (unconditional)

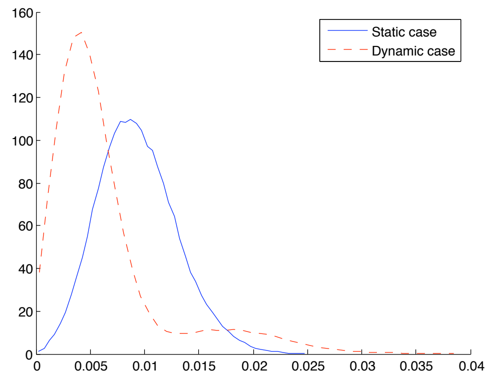

This figure contrasts the simulated loss density (pdf), of the basic model (no industry and seniority), for the static version (i.e. default probability and recovery rate distribution do not vary with the credit cycle) and the dynamic version (default probability and recovery rate distribution depend on the credit cycle). For the dynamic model, the probability of being in a credit downturn is assumed to be equal 0.

Figure 12: Simulated loss density (upturn)

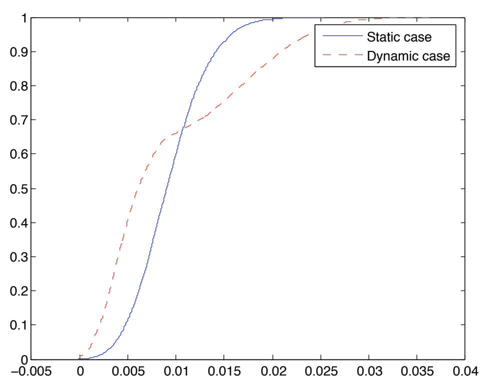

This figure contrasts the simulated loss distribution (cdf), of the basic model (no industry and seniority), for the static version (i.e. default probability and recovery rate distribution do not vary with the credit cycle) and the dynamic version (default probability and recovery rate distribution depend on the credit cycle). For the dynamic model, the probability of being in a credit downturn is assumed to be equal to the unconditional probability

Figure 13: Simulated loss distribution

CEMFI WORKING PAPERS

0401 *Andres Almazan, Javier Suarez and Sheridan Titman*: "Stakeholders, transparency and capital structure".

0402 *Antonio Diez de los Rios*: "Exchange rate regimes, globalisation and the cost of capital in emerging markets".

0403 *Juan J. Dolado and Vanessa Llorens*: "Gender wage gaps by education in Spain: Glass floors vs. glass ceilings".

0404 *Sascha O. Becker, Samuel Bentolila, Ana Fernandes and Andrea Ichino*: "Job insecurity and children's emancipation".

0405 *Claudio Michelacci and David Lopez-Salido*: "Technology shocks and job flows".

0406 *Samuel Bentolila, Claudio Michelacci and Javier Suarez*: "Social contacts and occupational choice".

0407 *David A. Marshall and Edward Simpson Prescott*: "State-contingent bank regulation with unobserved actions and unobserved characteristics".

0408 *Ana Fernandes*: "Knowledge, technology adoption and financial innovation".

0409 *Enrique Sentana, Giorgio Calzolari and Gabriele Fiorentini*: "Indirect estimation of conditionally heteroskedastic factor models".

0410 *Francisco Peñaranda and Enrique Sentana*: "Spanning tests in return and stochastic discount factor mean-variance frontiers: A unifying approach".

0411 *F. Javier Mencía and Enrique Sentana*: "Estimation and testing of dynamic models with generalised hyperbolic innovations".

0412 *Edward Simpson Prescott*: "Auditing and bank capital regulation".

0413 *Víctor Aguirregabiria and Pedro Mira*: "Sequential estimation of dynamic discrete games".

0414 *Kai-Uwe Kühn and Matilde Machado*: "Bilateral market power and vertical integration in the Spanish electricity spot market".

0415 *Guillermo Caruana, Liran Einav and Daniel Quint*: "Multilateral bargaining with concession costs".

0416 *David S. Evans and A. Jorge Padilla*: "Excessive prices: Using economics to define administrable legal rules".

0417 *David S. Evans and A. Jorge Padilla*: "Designing antitrust rules for assessing unilateral practices: A neo-Chicago approach".

0418 *Rafael Repullo*: "Policies for banking crises: A theoretical framework".

0419 *Francisco Peñaranda*: "Are vector autoregressions an accurate model for dynamic asset allocation?"

0420 *Ángel León and Diego Piñeiro*: "Valuation of a biotech company: A real options approach".

0421 *Javier Alvarez and Manuel Arellano*: "Robust likelihood estimation of dynamic panel data models".

0422 *Abel Elizalde and Rafael Repullo*: "Economic and regulatory capital. What is the difference?".

0501 C*laudio Michelacci and Vincenzo Quadrini*: "Borrowing from employees: Wage dynamics with financial constraints".

0502 *Gerard Llobet and Javier Suarez*: "Financing and the protection of innovators".

0503 *Juan-José Ganuza and José S. Penalva*: "On information and competition in private value auctions".

0504 *Rafael Repullo*: "Liquidity, risk-taking, and the lender of last resort".

0505 *Marta González and Josep Pijoan-Mas*: "The flat tax reform: A general equilibrium evaluation for Spain".

0506 *Claudio Michelacci and Olmo Silva*: "Why so many local entrepreneurs?".

0507 *Manuel Arellano and Jinyong Hahn*: "Understanding bias in nonlinear panel models: Some recent developments".

0508 *Aleix Calveras, Juan-José Ganuza and Gerard Llobet*: "Regulation and opportunism: How much activism do we need?".

0509 *Ángel León, Javier Mencía and Enrique Sentana*: "Parametric properties of semi-nonparametric distributions, with applications to option valuation".

0601 *Beatriz Domínguez, Juan José Ganuza and Gerard Llobet*: "R&D in the pharmaceutical industry: a world of small innovations".

0602 *Guillermo Caruana and Liran Einav*: "Production targets".

0603 *Jose Ceron and Javier Suarez*: "Hot and cold housing markets: International evidence".

0604 *Gerard Llobet and Michael Manove*: "Network size and network capture".

0605 *Abel Elizalde*: "Credit risk models I: Default correlation in intensity models".

0606 *Abel Elizalde*: "Credit risk models II: Structural models".

0607 *Abel Elizalde*: "Credit risk models III: Reconciliation reduced – structural models".

0608 *Abel Elizalde*: "Credit risk models IV: Understanding and pricing CDOs".

0609 *Gema Zamarro*: "Accounting for heterogeneous returns in sequential schooling decisions".

0610 *Max Bruche*: "Estimating structural models of corporate bond prices".

0611 *Javier Díaz-Giménez and Josep Pijoan-Mas*: "Flat tax reforms in the U.S.: A boon for the income poor".

0612 *Max Bruche and Carlos González-Aguado*: "Recovery rates, default probabilities and the credit cycle".